THE STRONG ATIYAH CONJECTURE AND VON NEUMANN RANKS FOR GRAPHS OF GROUPS WITH FINITE-INDEX EDGE INCLUSIONS

NIDHAL MGHIRBI AND SETH DOUGLAS

Abstract. Let G be the fundamental group of a graph of groups in which each edge group has finite index in both incident vertex groups. If the von Neumann matrix ranks at each vertex generate a discrete subgroup of the real numbers, we prove that the rank group of G is exactly the sum of the vertex rank groups. This gives permanence for the strong Atiyah conjecture with bounded torsion. The proof expresses a rank as a minimum over edge projections and shows that modules over the edge regular closures attain it simultaneously. In the torsion-free case with integral vertex ranks, the affiliated division ring is the universal field of the graph of vertex division rings. For finite graphs of finitely generated free groups, rational matrix ranks are computable. Applications include finitely presented infinite simple groups acting on products of trees and the exact conjectured denominator for every subgroup of $\mathrm{GL}_2(\mathbb{Q})$. In $\mathrm{SL}_2(\mathbb{Z}[1/2])$ an integral group-ring element has kernel dimension $1/12$, although every torsion-free subgroup of finite index has index divisible by 24.

Claude Opus 5.5, GPT-6.1 Sol, GPT-6.2 Sol and GPT-6.2 Astra assisted with this work; the authors take responsibility for every claim.

1. Introduction

Let K be a subfield of $\mathbb{C}$ closed under complex conjugation. Write $\Lambda_K(H)$ for the additive subgroup of $\mathbb{R}$ generated by the von Neumann ranks of all finite rectangular matrices over $K[H]$. We take the sum of the diagonal traces, so that the identity matrix of size n has rank n. Put

$$d(H) = \mathrm{lcm}\{|F| : F \le H \text{ finite}\}.$$

When $d(H) < \infty$, the strong Atiyah conjecture over K says that $\Lambda_K(H) \subseteq d(H)^{-1}\mathbb{Z}$. Equality then follows from the projections $|F|^{-1}\sum_{g\in F} g$, of rank $1/|F|$.

Theorem A. *Let G be the fundamental group of a connected graph of groups. Suppose that every edge group has finite index in each of its incident vertex groups and that $\Lambda_K(G_v)$ is discrete for each vertex. Then*

$$\Lambda_K(G) = \sum_v \Lambda_K(G_v).$$

In particular, if $d(G) < \infty$ and all vertex groups satisfy the strong Atiyah conjecture over K, then so does G.

The inclusions can be nonnormal, the two indices can differ, and several edges can meet at a vertex. Every finite subgroup fixes a vertex of the Bass–Serre tree. Thus the denominators supplied by the vertices in the last assertion are exactly the ones the conjecture predicts. The right side in the first assertion means finite sums, also for an infinite graph.

Theorem B. *For every subgroup $H \le \mathrm{GL}_2(\mathbb{Q})$ one has $\Lambda_{\mathbb{C}}(H) = d(H)^{-1}\mathbb{Z}$. The denominators of $\mathrm{GL}_2(\mathbb{Q})$ and $\mathrm{SL}_2(\mathbb{Q})$ are 24 and 12, respectively. For $H = \mathrm{SL}_2(\mathbb{Z}[1/N])$, $N \ge 1$, the denominator is 12. A single element of $\mathbb{Z}[\mathrm{SL}_2(\mathbb{Z})]$ has kernel dimension $1/12$ in every $\mathrm{SL}_2(\mathbb{Z}[1/N])$ and in $\mathrm{SL}_2(\mathbb{Q})$.*

Date: Version 2.0.0, November 2026.

For $N = 2$, a torsion-free subgroup of index 24 gives a denominator bound of 24 by restriction of scalars. Every torsion-free subgroup of finite index has index divisible by 24 (Proposition 7.4). Theorem B gives the sharp denominator 12, and one integral group-ring element detects it.

How ranks on the two sides fit together. Consider first one edge E joining groups V_0, V_1, and a product XY in which X comes from V_0 and Y from V_1. A projection w over the edge selects part of the intermediate space. Splitting the product there gives

$$XY = XwY + X(1-w)Y, \qquad \mathrm{rk}_G(XY) \le \mathrm{rk}_{V_0}(Xw) + \mathrm{rk}_{V_1}((1-w)Y).$$

The useful point is that one can choose w so that equality holds. Section 2 derives this from the atom theorem of Belinschi, Bercovici and Liu. Appendix A gives a second proof for a single group amalgam by rotating two copies of its regular representation. For graphs with two or more edges the formula rests on the atom theorem alone.

The first difficulty already occurs on a two-edge path $V_0 *_{E_1} V_1 *_{E_2} V_2$. After forming the first amalgam, E_2 can have infinite index in it. We therefore cut at both edges at once. In a general graph we gather the rows and columns at each vertex into one block, and put all edge algebras on one diagonal.

At each vertex we look for a submodule having a large intersection with a given kernel or image, charging a price t for each unit of dimension. Finite-index coset coordinates show that the greatest maximizing submodule is the same whichever incident edge we use to find it. Taking its identity-coset coefficient supplies a submodule at each edge. Integrating the vertex dimensions from $t = 0$ to $t = 1$ shows that the costs add to the desired rank for almost every common choice of t. These costs are dimensions of vertex projectives, which retain their actual von Neumann dimensions even when the vertices have torsion.

Sánchez-Peralta proved permanence for countable groups with bounded torsion and finite edge groups [21, Theorem 1.1]. His Theorem 1.6 covers the countable groups with bounded torsion in his class $\mathcal{T}_{\mathrm{VLI}}$ [21, Definition 1.5 and the paragraph after it], which includes graphs of virtually free groups whose fundamental group is virtually locally indicable. Theorem A allows an infinite edge group of finite index at each end, with no local indicability hypothesis. When the vertex ranks are integers, Fisher and Sánchez-Peralta had already obtained an abstract division-ring embedding for arbitrary edge groups [7, Corollary 3.12]. In the finite-index case we identify the universal field of the graph of division rings *inside the regular representation* (Theorem 6.2). Its matrix ranks can consequently be computed by factoring matrices. For a finite graph of finitely generated free groups this gives an exact algorithm (Theorem 6.4).

When only rational coefficients are understood at the vertices, it suffices to know integral ranks on all their finite direct powers (Corollary 6.1).

2. A rank formula over a common von Neumann algebra

For a finite von Neumann algebra M with faithful normal tracial state τ, write $U(M)$ for its affiliated ring. For a group H we abbreviate $U(\mathcal{N}(H)) = \mathcal{U}(H)$ and $U_H = \mathcal{U}(H)$. We take the left regular convention, with $\tau_H(a) = \langle a\delta_1, \delta_1\rangle$. On matrices Tr_M is the sum of the diagonal traces. In particular $\mathrm{rk}_M I_n = n$. Background on affiliated operators and dimension can be found in Berberian [2], Reich [20] and Lück [14, Chapters 6 and 8].

We use right modules and let matrices act on columns on the left. Polar decomposition and inversion of the absolute value on its support give a partial inverse $T^\dagger$ to an affiliated matrix. Consequently

$$\mathrm{im}\, T = (TT^\dagger)U(M)^a, \qquad \ker T = (1 - T^\dagger T)U(M)^b \quad (T \text{ of size } a \times b).$$

Both generating matrices are orthogonal projections. The dimension of $pU(M)^a$ is $\mathrm{Tr}_M(p)$, and $\mathrm{rk}_M T$ is the dimension of its image. For bounded T these are the Hilbert-space kernel and closed-range dimensions. A square matrix of full rank is invertible in the affiliated ring.

Finite sums, intersections and quotients of projection-generated submodules are finitely generated projective. Dimension is additive and detects zero among these modules. For instance, $pU(M)^a + qU(M)^a = (p \vee q)U(M)^a$, by taking the range of the row $(p\ q)$. For an infinite join we take the module of the supremum of the projections. Inclusion of a subgroup preserves affiliated matrix supports and their traces, since the regular representation restricts to a sum of copies of the subgroup representation.

Let M_0, M_1 be unital von Neumann subalgebras of M containing $\mathcal{B}$, with its trace-preserving conditional expectation $\mathbb{E} : M \to \mathcal{B}$. Freeness over $\mathcal{B}$ means that $\mathbb{E}(x_1 \cdots x_k) = 0$ whenever the x_j come alternately from the two algebras and $\mathbb{E}(x_j) = 0$ for all j. This property passes to matrix amplifications, by expanding a matrix entry. For a group amalgam $A *_E B$, normal form gives freeness of $\mathcal{N}(A)$ and $\mathcal{N}(B)$ over $\mathcal{N}(E)$. Indeed a reduced word in elements outside the edge has zero edge coefficient. Centered group-ring approximations, bounded in operator norm and converging in L^2, give the assertion for the von Neumann algebras.

Theorem 2.1 (Analytic edge formula)**.** *Suppose M_0, M_1 are free over $\mathcal{B}$ as above. For $X \in M_{a,m}(U(M_0))$ and $Y \in M_{m,b}(U(M_1))$,*

$$\mathrm{rk}_M(XY) = \min_{\substack{w=w^*=w^2 \\ w\in M_m(\mathcal{B})}} \{\, \mathrm{rk}_{M_0}(Xw) + \mathrm{rk}_{M_1}((1-w)Y)\}. \tag{2.1}$$

If p, q are the kernel projection of X and range projection of Y, the least minimizing projection is $\operatorname{supp} \mathbb{E}_m(p \wedge q)$*, where $\mathbb{E}_m$ acts entrywise.*

Proof. Every edge projection gives an upper bound by writing $XY = XwY + X(1-w)Y$. Set $h = p \wedge q$ and $e = \operatorname{supp} \mathbb{E}_m(h)$. Faithfulness applied to $(1-e)h(1-e)$ gives $h \le e$. Conversely an edge projection containing h contains e. For any edge projection w, restrict X to the range of w and $1-w$ to the range of Y. This gives

$$\begin{aligned} \mathrm{rk}_{M_0}(Xw) &= \mathrm{Tr}(w) - \mathrm{Tr}(p \wedge w), \\ \mathrm{rk}_{M_1}((1-w)Y) &= \mathrm{Tr}(q) - \mathrm{Tr}(q \wedge w), \\ \mathrm{rk}_M(XY) &= \mathrm{Tr}(q) - \mathrm{Tr}(h). \end{aligned}$$

Thus it suffices to prove

$$\mathrm{Tr}(h) = \mathrm{Tr}(p \wedge e) + \mathrm{Tr}(q \wedge e) - \mathrm{Tr}(e). \tag{2.2}$$

The case $h = 0$ gives $e = 0$. Otherwise work in the corner of e, normalizing the trace there. Put $p' = p \wedge e,\ q' = q \wedge e$. Their meet is h, and their compressed vertex algebras remain free over $eM_m(\mathcal{B})e$. Indeed the expectation of a corner element already belongs to the corner.

First suppose $\mathbb{E}_m(h)$ is invertible in the corner. Positivity of $e - p'$ and $e - q'$ identifies h with the kernel projection of $2e - p' - q'$. Apply the atom theorem of Belinschi, Bercovici and Liu at $2e$ [1, Theorem 3.5(vii), arXiv v2]. It supplies selfadjoints b_0, b_1 in the edge corner, with sum $2e$, whose kernel projections for $b_0 - p'$ and $b_1 - q'$ have meet h and traces adding to $1 + \tau(h)$ in the normalized corner. Both $b_i h = h$. Taking expectation gives $b_i \mathbb{E}_m(h) = \mathbb{E}_m(h)$, whence $b_i = e$. The kernel projections are therefore p', q', proving (2.2) upon restoring the trace.

In general let $d_n = 1_{[0,1/n)}(\mathbb{E}_m(h))$ in the corner of e. Its trace tends to zero because the support of $\mathbb{E}_m(h)$ is e. Enlarge p', q' to $p_n = p' \vee d_n,\ q_n = q' \vee d_n$ in their own vertex algebras. Their meet h_n dominates both h and d_n, so

$$\mathbb{E}_m(h_n) \ \ge\ \tfrac{1}{2}(\mathbb{E}_m(h) + d_n) \ \ge\ (2n)^{-1} e.$$

The invertible case applies to p_n, q_n. Each enlargement costs at most $\operatorname{Tr}(d_n)$, giving

$$\operatorname{Tr}(h) \le \operatorname{Tr}(h_n) = \operatorname{Tr}(p_n) + \operatorname{Tr}(q_n) - \operatorname{Tr}(e) \le \operatorname{Tr}(p') + \operatorname{Tr}(q') - \operatorname{Tr}(e) + 2\operatorname{Tr}(d_n).$$

The reverse inequality follows from $p' \vee q' \le e$. Passing to the limit proves (2.2).

Finally, if w attains the minimum, the three cost identities give

$$\operatorname{Tr}(h) = \operatorname{Tr}(p \wedge w) + \operatorname{Tr}(q \wedge w) - \operatorname{Tr}(w) \le \operatorname{Tr}(h \wedge w) \le \operatorname{Tr}(h).$$

Faithfulness forces $h \le w$, hence $e \le w$. This proves leastness. □

Example. For a scalar edge and one intermediate coordinate the only edge projections are 0 and 1, so (2.1) gives the smaller of the two ranks. In $G = \langle s, t \mid s^2 = t^2 = 1\rangle$, take $X = (1+s)/2$, $Y = (1+t)/2$. Each has trace $1/2$. A vector in $\ker X \cap \operatorname{im} Y$ satisfies $st\xi = -\xi$, so its absolute value is constant on each orbit of st. The orbits are infinite, and square summability forces $\xi = 0$. Thus $\operatorname{rk}_G(XY) = 1/2$. With two intermediate coordinates an actual line can be needed. For $X = (1,1)$ and $Y = (1,-1)^{\mathsf{T}}$, both costs vanish for the projection onto $\operatorname{im} Y$.

With trivial edge group, the formula also explains the torsion-free free-product case of Sánchez-Peralta's permanence theorem [21, Theorem 1.1]. Suppose the vertex groups have integral ranks over $\mathbb{C}$. Their affiliated division closures are division rings by Lemma 3.1 below. Every minimizing projection in (2.1) has complex entries, so the two costs for vertex matrices can be calculated by elimination over the vertex division rings. They are integers. An arbitrary group-ring matrix can be put into this product form after adding identity blocks, by the elementary stabilization in the proof of Theorem A. This gives integrality for the free product, even when the trivial subgroup has infinite index at both ends.

There is a field statement here too. In the notation of Section 6, the affiliated division closure of the free product is the universal field of the coproduct $R_{V_0} *_{\mathbb{C}} R_{V_1}$; this is the torsion-free case of Sánchez-Peralta's universal localization theorem [21, Theorems 1.2(iii) and 3.8]. Indeed, splitting at the minimizing complex subspace bounds the rank in every division-ring image of this coproduct by the same two costs. Stabilization gives the comparison for every matrix, and the universal-field argument in Theorem 6.2 applies verbatim. In this case existence of a subspace over the coefficient field follows directly from the analytic formula.

3. Regular closures and restriction of scalars

Let R_H be the $*$-regular closure of $K[H]$ in U_H. Concretely we close under ring operations, adjoints and the entries of Moore–Penrose inverses of affiliated matrices. This is also the usual $*$-regular closure: a $*$-regular subring of U_H is regular in all matrix sizes, and the unique partial inverse with orthogonal initial and final supports is computed there. One may use the concrete closure throughout the proof.

For a finitely generated submodule L of R_H^n put

$$\delta_H(L) = \dim_H(LU_H).$$

The tensor extension here embeds in U_H^n, since an image over a regular ring is a direct summand. The same prescription gives a faithful additive dimension on finitely generated projectives. If $J \le H$, restriction of the regular representation of H to J is a sum of copies of the regular representation of J. Thus $U_J \hookrightarrow U_H$ preserves matrix partial inverses and normalized traces, $R_J \hookrightarrow R_H$, and $\Lambda_K(J) \subseteq \Lambda_K(H)$.

Lemma 3.1. *Matrix ranks over R_H belong to $\Lambda_K(H)$. If this subgroup is discrete, R_H is semisimple Artinian. The same conclusion holds for R_J for every subgroup $J \le H$.*

Proof. The first assertion also follows from Jaikin-Zapirain [9, Corollary 6.2]. We give a dimension argument to specify what closure is needed. Regard the entries of a column

as variables in the right module U_H. The graph of a matrix over $K[H]$ is defined by homogeneous linear equations with these coefficients. Relations described by a finite system of this kind together with existentially quantified auxiliary columns are closed under sums of maps, compositions, direct sums and coordinate projections. Graphs of matrices assembled from such maps have the same description.

In passing to the partial inverse of A use the equations

$$z = A^* w, \qquad A^*(Az - y) = 0.$$

For given y the unique z solving these equations for some w is $A^\dagger y$. Indeed Az is the component of y in $\operatorname{im} A$, and A is an isomorphism from $\operatorname{im} A^*$ onto $\operatorname{im} A$. Maintain descriptions for the graphs of each expression and its adjoint simultaneously. The adjoint reverses products, and $(A^\dagger)^* = (A^*)^\dagger$. This proves by induction that the image of every matrix over R_H is the coordinate projection of the kernel of a matrix over $K[H]$. The dimension of this projection is the dimension of the full kernel minus the dimension of its submodule where the retained coordinates vanish. Both are integers minus group-ring matrix ranks, as desired.

A discrete subgroup containing 1 is $D^{-1}\mathbb{Z}$ for a positive integer D. A nonzero projection in R_H has trace at least $1/D$. If pR_H contains a nonzero proper right ideal, choose nonzero x in it. Then $xx^\dagger R_H = xR_H$ is a nonzero proper summand of pR_H, and $xx^\dagger \le p$. Splitting orthogonal projections in this way must end in at most D pieces beneath 1. This writes the regular module as a finite sum of simple modules. Subgroup ranks lie in the same grid as ambient ranks, which proves the last assertion without any hypothesis about $d(J)$. □

Lemma 3.2 (Coset coordinates)**.** *Suppose $E \le V$ has finite index s. Then*

$$R_V = \bigoplus_{rE \in V/E} rR_E, \qquad U_V = \bigoplus_{rE \in V/E} rU_E, \qquad R_V \otimes_{R_E} U_E \simeq U_V.$$

There are also left-module coordinates. In either decomposition the identity-coset coefficient $\varepsilon_{V,E}$ is the same U_E-bimodule map on U_V, and it takes R_V into R_E. Restriction of a projection-generated U_V-module P satisfies $\dim_E P = s \dim_V P$.

Proof. Use the unitary map from s copies of $\ell^2(E)$ to $\ell^2(V)$ that sends a tuple (ξ_r) to $\sum_r r\xi_r$. Left convolution gives a unital embedding $U_V \hookrightarrow M_s(U_E)$. For bounded operators its entries are subgroup expectations of $r^{-1}xs'$ for representatives r, s'. The embedding extends to affiliated operators by polar decomposition and inversion on support. Its image of $K[V]$ has entries in $K[E]$. Since the embedding preserves adjoints and partial inverses, the image of R_V has entries in R_E.

For any $x \in U_V$ its column at the identity coset has entries $u_r \in U_E$. The element $\sum_r ru_r$ has the same column. An affiliated operator with zero identity column vanishes, since its kernel projection contains the whole identity-coset Hilbert space and commutes with right translations by V. Hence the columns determine the operator, giving the two right-module decompositions. The case $x \in R_V$ uses $u_r \in R_E$ and $R_E \subset R_V$. Adjoints give left coordinates.

To compare the two coefficient maps, first use the finite-index normal core $N = \bigcap_{v \in V} vEv^{-1}$. Conjugation preserves U_N, so $U_V = \bigoplus_{rN} rU_N$ is graded by V/N, with each $rU_N = U_N r$. The projection keeping precisely the grades in E/N is the identity-coset coefficient for both E-coordinate systems. It is U_E-bilinear. On bounded operators this coefficient is the trace-preserving expectation onto $\mathcal{N}(E)$.

For a matrix projection p generating P, its diagonal coset blocks are the identity coefficients of $r^{-1}pr$. Each has E-trace equal to the V-trace of p. Summing over r gives the dimension formula. □

4. A maximizing module independent of the incident edge

Suppose $\Lambda_K(V)$ is discrete, let $E \le V$ have finite index, and let $S \subseteq U_V^n$ be projection-generated over U_V. For a right R_E-submodule $L \subset R_V^n$, write $L_U = LU_E$. These are finite-length rational modules by Lemmas 3.1 and 3.2. Exact sequences split, so tensor extension commutes with intersections, finite sums and preimages. In particular $\delta_E(L) = \dim_E L_U$, also in the coset-coordinate module.

Lemma 4.1. *For $0 \le t \le 1$ there is a greatest submodule maximizing*

$$h_{E,t}(L) = \dim_E(L_U \cap S) - t\delta_E(L).$$

It is a right R_V-submodule $P_t(S)$ independent of E. Furthermore

$$\int_0^1 \delta_V(P_t(S))\,dt = \dim_V S.$$

Proof. Dimension additivity gives

$$h_{E,t}(L + L') + h_{E,t}(L \cap L') \;\ge\; h_{E,t}(L) + h_{E,t}(L'),$$

because the sum of the two overlaps lies in the overlap of the sum. The rational dimensions $\delta_E(L)$ take values in a finite set: they are in the vertex grid by Lemma 3.1 and bounded by $n[V : E]$. A proper inclusion increases dimension strictly. Write $\sigma = \sup_L h_{E,t}(L)$. Choose the largest dimension j at which the supremum over that dimension equals σ. If larger dimensions occur, their scores are at most $\sigma - \eta$ for some $\eta > 0$. Two distinct modules of dimension j with scores greater than $\sigma - \eta/3$ would have a sum of larger dimension with score greater than $\sigma - 2\eta/3$. Thus sufficiently close approximants coincide, giving attainment. If there is no larger dimension the only module of that dimension is the ambient module. Sums of maximizers maximize, giving a greatest maximizer in all cases.

Use the normal core N first. Right multiplication of columns by $g \in V$ is semilinear for $a \mapsto g^{-1}ag$ on R_N, U_N. It preserves the N-dimensions of both L_U and $L_U \cap S$: the twist is trace preserving, its coordinate change is invertible, and $Sg = S$. Uniqueness makes the greatest N-maximizer invariant under every g. It is an R_V-module by the normal coset decomposition. For an R_E-module, $LU_E = LU_N$ because $U_E = R_E U_N$. Its N-score is $[E : N]$ times its E-score. As the greatest N-maximizer is already an R_E-module it is also the greatest E-maximizer. On R_V-modules the E-score equals $[V : E]$ times the V-score. This proves independence, even for different incident subgroups.

Set $E = V$. The maximal score as a function of t is the upper envelope of finitely many lines, one for each possible rational dimension, with intercept the supremum of the overlap in that dimension. Off the finitely many breakpoints its slope is $-\delta_V(P_t(S))$. At 0 the whole module gives $\dim_V S$. At 1 the maximum is zero since an overlap has dimension at most that of the rational module extended to U_V. Integration proves the formula. □

Lemma 4.2. *Suppose E embeds with finite index in V_1, V_0 as above. Let $S \subset U_{V_1}^b$, $T \subset U_{V_0}^a$ be projection-generated vertex modules. If an R_E-linear map $c : R_{V_1}^b \to R_{V_0}^a$ extends in coset coordinates to a map carrying S into T, then*

$$c(P_t(S)) \subseteq P_t(T).$$

Proof. Write $Q = P_t(S), P = P_t(T), J = Q \cap c^{-1}P, P' = P + cQ$. The isomorphic R_E-modules Q/J and P'/P have the same weight k. If $k > 0$, P' strictly contains the greatest maximizer. Writing $d_S(L) = \dim_E(L_U \cap S)$, and similarly for T, maximality gives

$$d_S(Q) - d_S(J) \;\ge tk \;>\; d_T(P') - d_T(P).$$

On the other hand c injects $(Q_U \cap S)/(J_U \cap S)$ into $(P'_U \cap T)/(P_U \cap T)$. Its kernel computation follows by tensoring the split exact kernel sequence for $Q \to R_{V_0}^a/P$. This contradicts the displayed inequality. Thus $k = 0$, and dimension faithfulness gives the inclusion. □

5. One module at each edge

For a single amalgam, Theorem 2.1 gives one minimizing edge projection. For the path $A *_{E_1} B *_{E_2} C$, put the two edges in the two coordinates of a column. Then the algebras on the two sides are

$$\mathcal{N}(A) \oplus \mathcal{N}(C) \qquad \text{and} \qquad M_2(\mathcal{N}(B)).$$

Their common diagonal is $\mathcal{N}(E_1) \oplus \mathcal{N}(E_2)$. A matrix entry at B may carry a vector from the coordinate for E_1 to the coordinate for E_2. This single off-diagonal entry is what is lost if the two gluings are treated successively.

5.1. **The matrix algebra of a finite graph.** Consider a finite connected bipartite graph of groups with at least one edge. Write V_0, V_1 for the two vertex classes, I for the set of geometric edges, and $d = |I|$. Orient e from its endpoint $v \in V_1$ to its endpoint $u \in V_0$. Let I_v be the set of incident edges at a vertex v, $d_v = |I_v|$, and $\iota_{v,e} : G_e \to G_v$ the incidence homomorphism. Parallel edges have separate coordinates. Choose a maximal tree. In the resulting presentation of G the stable letters satisfy

$$t_e \iota_{v,e}(g) t_e^{-1} = \iota_{u,e}(g) \quad (g \in G_e), \qquad t_e = 1 \quad (e \text{ in the tree}).$$

Every graph can be made bipartite by inserting at the midpoint of each edge a vertex with that edge group and identity inclusions on its side. For a loop this gives two distinct parallel edges. Subdivision preserves the fundamental group and the finite-index hypothesis.

Use I as the matrix coordinates and write $J = \operatorname{diag}(t_e)_{e \in I}$. Inside $M_d(\mathcal{N}(G))$, with its trace Tr_G / d, set

$$M_0 = \bigoplus_{u \in V_0} M_{d_u}(\mathcal{N}(G_u)), \qquad M_1 = J \left(\bigoplus_{v \in V_1} M_{d_v}(\mathcal{N}(G_v)) \right) J^*. \tag{5.1}$$

A block for v occupies the positions $I_v \times I_v$. The partitions of I used on the two sides are generally different. The diagonal copy of $\mathcal{B} = \bigoplus_{e \in I} \mathcal{N}(G_e)$ uses the left endpoint embedding at each e, and the stable-letter relations give the same copy from the right. Its trace-preserving conditional expectation discards off-diagonal entries and takes the subgroup expectation at each diagonal entry.

Lemma 5.1 (The two vertex algebras)**.** *The algebras in* (5.1) *are free over* $\mathcal{B}$*. The corresponding two block algebras over the vertex group rings generate* $M_d(K[G])$*.*

Proof. For freeness, regard the matrix coordinates as points in the interiors of edges. An elementary block entry from f to e through $u \in V_0$ is gE_{ef}, with $g \in G_u$ and E_{ef} a matrix unit. Through $v \in V_1$ it is $t_e g t_f^{-1} E_{ef}$, $g \in G_v$. Both are the group elements read on the corresponding paths, after closing them by the chosen tree paths. More explicitly, on further subdivision at these interior points take a maximal tree containing every half-edge on the V_0 side, and both halves of every original tree edge. The half-edge left over at a nontree edge then carries exactly t_e, in the orientation used above.

An off-diagonal block entry passes through a vertex by two different edges. A centered diagonal entry is a sum of entries which return by the same edge but have intervening group label outside its incidence subgroup. In either case the path is reduced at that vertex. In an alternating product the two passages meeting at an interior point use opposite ends of its edge. There is no cancellation there either. Noncomposable entries multiply to zero. A remaining diagonal entry is a reduced path of positive length based at the interior point, and Bass–Serre normal form prevents its value from lying in the group G_e of that point [22, Chapter I, §5]. Its edge expectation is zero. This proves freeness for the group algebras. Bounded L^2-approximation and subtraction of the expectation extend it to M_0, M_1 because products of uniformly bounded approximants converge in L^2 by telescoping.

For generation, two tree edges with a common endpoint are joined by ordinary matrix units on the appropriate side, since their letters are 1. Tree edges form a connected set under such moves. Every nontree edge is joined by ordinary matrix units to a tree edge at its V_0 endpoint. Thus all ordinary matrix units are generated. Every vertex group occurs untwisted in a diagonal coordinate at an incident tree edge. For a nontree edge e, choose a tree edge f at its V_1 endpoint. The right-hand block contains $t_e E_{ef}$ and $t_e^{-1} E_{fe}$. Transporting these entries by ordinary matrix units gives the stable letters and their inverses on each diagonal coordinate. These elements generate $M_d(K[G])$. □

5.2. **Attaining the minimum over the regular closures.** For a product XY with X in the left block algebra over R_{G_u} and Y in the right one over R_{G_v}, write X_u, Y_v for the local matrices before inserting the stable letters. If the unexpanded sizes are $a \times m$ and $m \times b$, the local sizes are $ad_u \times md_u$ and $md_v \times bd_v$. An edge has m intermediate coordinates. For $L_e \subset R_{G_e}^m$ a right submodule, put

$$L(v) = \bigoplus_{e \in I_v} \iota_{v,e}(L_e) R_{G_v}, \qquad \pi_v : R_{G_v}^{md_v} \longrightarrow R_{G_v}^{md_v}/L(v).$$

Proposition 5.2. *With the unnormalized trace of the expanded group matrix,*

$$\mathrm{rk}_G(XY) = \min_{(L_e)_e} \left\{ \sum_{u \in V_0} \delta_{G_u}(X_u L(u)) + \sum_{v \in V_1} \delta_{G_v}(\mathrm{im}(\pi_v Y_v)) \right\}.$$

Proof. Each L_e is generated by a matrix projection over R_{G_e}, obtained by taking the range projection of a matrix of generators. Extension of its split image and quotient computes the two costs over U_{G_v}. Consequently every rational tuple gives an upper bound by the analytic edge formula, with its trace multiplied by $d = |I|$.

Take an analytic minimizing tuple (w_e). In the local coordinates put

$$w(v) = \bigoplus_{e \in I_v} \iota_{v,e}(w_e), \qquad S_v = \ker((1 - w(v))Y_v), \qquad T_u = \mathrm{im}(X_u w(u)).$$

The rank equals $\sum_u \dim_u T_u + \sum_v (bd_v - \dim_v S_v)$, where $\dim_v$ denotes affiliated dimension at G_v. Put $Q_v = P_t(S_v)$, $P_u = P_t(T_u)$. On an edge $e : v \to u$ use the block of rows $Y_{v,e}$ and the block of columns $X_{u,e}$. Define the R_{G_e}-linear map

$$c_e(x) = X_{u,e}\, \varepsilon_{G_v, G_e}(Y_{v,e} x),$$

transporting the edge coefficient from its incidence at v to its incidence at u. For $x \in S_v$, $Y_{v,e} x = w_e Y_{v,e} x$. Bimodularity of the coefficient map gives

$$\varepsilon_{G_v, G_e}(Y_{v,e} x) = w_e\, \varepsilon_{G_v, G_e}(Y_{v,e} x).$$

Thus the extension of c_e maps S_v into T_u, and Lemma 4.2 gives $c_e(Q_v) \subseteq P_u$.

Set $L_e(t) = \varepsilon_{G_v, G_e}(Y_{v,e} Q_v)$. At u we have $X_u L(u) \subseteq P_u$. At v we need all coset coefficients. With representatives for $G_e \backslash G_v$, write $Y_{v,e} x = \sum_r z_r r$. Right linearity of $Y_{v,e}$ gives

$$z_r = \varepsilon_{G_v, G_e}\big(Y_{v,e}(x r^{-1})\big).$$

For $x \in Q_v$ the element $x r^{-1}$ belongs to the very same vertex module Q_v. Therefore each $z_r \in L_e(t)$, and $Y_v Q_v \subseteq L(v)$. If $C(t)$ denotes the proposed cost, the two inclusions give

$$\mathrm{rk}_G(XY) \le C(t) \le \sum_u \delta_{G_u}(P_u) + \sum_v \big(bd_v - \delta_{G_v}(Q_v)\big).$$

By Lemma 4.1 the right side integrates to the analytic rank, and it is a step function up to finitely many breakpoints. Equality holds almost everywhere. Any one such t attains the minimum simultaneously at all edges. □

Proof of Theorem A. Embed a rectangular matrix M over $K[G]$ in one matrix corner of $M_d(K[G])$ by replacing each entry x by xE_{ff} for a fixed edge f. Its expanded rank is unchanged. The two vertex block rings generate $M_d(K[G])$ by Lemma 5.1. To separate their letters, split one word in an entry as a product and write the matrix as $B + UV$ with U a column and V a row. The operation

$$\begin{pmatrix} 1 & -U \\ 0 & 1 \end{pmatrix} \begin{pmatrix} B & U \\ -V & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ V & 1 \end{pmatrix} = \begin{pmatrix} B+UV & 0 \\ 0 & 1 \end{pmatrix}$$

removes one multiplication from the word expressions. Iteration gives a stabilized sum $A_0 + A_1 = [\, A_0 \; I \,]\binom{I}{A_1}$ with one factor on each side. Every stabilized identity contributes d to the expanded rank. By Proposition 5.2 and Lemma 3.1, the rank of M belongs to the sum of the local rank groups. Midpoint groups introduced in subdivision are subgroups of original vertex groups, so their ranks already occur there.

Conversely a matrix over a vertex subgroup has the same rank in G. For an infinite graph, finitely many entries use finitely many group words. Include the paths in a chosen maximal tree needed to connect their vertices and edges to a base point. The resulting finite connected subgraph has its fundamental group embedded in G by normal form and computes the same ranks. A one-vertex graph without edges is immediate. This proves the equality in all cases. For the assertion about denominators, finite vertex subgroups embed in G, and every finite subgroup of G fixes a vertex in its action without inversions on the Bass–Serre tree. □

5.3. **A consequence for projective modules.** The same attained tuple records more than the scalar rank. Write $K_0(R)$ for the group completion of the monoid of finitely generated projective right R-modules.

Corollary 5.3. *Under the hypotheses of Theorem A, if the sum of the vertex rank groups is discrete, extension of scalars gives a surjection*

$$\bigoplus_v K_0(R_{G_v}) \longrightarrow K_0(R_G).$$

For finite edge groups the corresponding surjection is due to Sánchez-Peralta [21, Theorem 1.2(iv)].

Proof. By Theorem A and Lemma 3.1 the target ring R_G is semisimple Artinian. Consider the stabilized product XY in the preceding proof. At an attaining edge tuple take the orthogonal rational projection a_e onto L_e and assemble them into a. The two maps

$$z \longmapsto (XaYz, (1-a)Yz), \qquad (\xi, \eta) \longmapsto \xi + X\eta$$

factor XY through $\operatorname{im}(Xa) \oplus \operatorname{im}((1-a)Y)$. This middle module is the sum of modules induced from the two sets of vertices, after conjugating the right blocks by the stable letters. Its dimension equals $\operatorname{rk}_G(XY)$. Dimension is faithful on finite projectives and all the images and kernels split, so the first map is surjective onto the middle module and the second is injective there. Thus $[\operatorname{im} XY]$ is in the image on K_0. Removing an identity block subtracts a multiple of $[R_G]$, itself induced. It follows that $[\operatorname{im} M]$ belongs to the image for every $K[G]$-matrix M.

Finally every finitely generated projective over R_G is the image of an idempotent e in a matrix ring. Jaikin-Zapirain's Corollary 6.2 [9], applied with a single coefficient matrix and identity multiplier, expresses the rank of e as a difference of ranks of two fixed $K[G]$-matrices, uniformly for every Sylvester matrix rank on R_G (pad by zeros for rectangular matrices). Apply this identity separately to the normalized ranks of the simple Artinian factors. These ranks together detect equality in $K_0(R_G)$: on each matrix ring over a division ring they record the multiplicity of its simple module up to a fixed nonzero normalization. Hence $[eR_G^n]$ is the difference of the two group-ring image classes. The finite-subgraph argument above suffices for these matrices even in an infinite graph. □

6. The affiliated division ring and exact ranks

Corollary 6.1 (Rational tests for complex coefficients)**.** *Let G be the fundamental group of a connected graph of groups with torsion-free vertex groups G_v and finite-index edge inclusions. Suppose that, for every vertex v and every $m \geq 1$, all matrices over $\mathbb{Q}[G_v^m]$ have integral von Neumann rank, where G_v^m is the direct product of m copies of G_v. Then G satisfies the strong Atiyah conjecture over $\mathbb{C}$.*

Proof. Since G_v is torsion-free, the companion coefficient theorem applies with G_v itself as the starting group [6, Theorem 1.2]. It gives $\Lambda_{\mathbb{C}}(G_v) = \mathbb{Z}$ for every vertex. Theorem A then gives $\Lambda_{\mathbb{C}}(G) = \mathbb{Z}$. The group G is torsion-free because every finite subgroup fixes a vertex in the Bass–Serre tree. □

The hypothesis for $m = 1$ alone gives the rational conclusion $\Lambda_{\mathbb{Q}}(G) = \mathbb{Z}$ directly by Theorem A. The passage to complex coefficients above uses the rational hypotheses on the higher powers of the vertex groups.

6.1. **The universal field in the torsion-free case.** Assume first that K is closed under complex conjugation and that $\Lambda_K(G_v) = \mathbb{Z}$ at every vertex. A nontrivial finite subgroup would give a projection of rank strictly between zero and one, so the vertex groups, and hence G, are torsion-free. Theorem A gives $\Lambda_K(G) = \mathbb{Z}$. By subgroup restriction and Lemma 3.1, the rings R_G, R_{G_v}, R_{G_e} are division rings: every nonzero scalar in one of them has affiliated rank one and its affiliated inverse belongs to the regular closure. They are also the division closures of the respective group rings. Indeed the division closure is then a division subring closed under adjoints, and therefore is itself $*$-regular.

Fix a maximal tree and let A_G be the fundamental ring of the graph of division rings R_{G_v}, R_{G_e}. It is generated by the vertex division rings and the stable letters, with the tree letters set to one and the edge relations imposed on all of R_{G_e}. Its canonical map to R_G sends the stable letters to the corresponding group elements. Fisher and Sánchez-Peralta's semifir construction applies to these division rings [7, Definition 3.1 and proof of Theorem 3.10]. Their Corollary 3.12 gives an abstract division-ring embedding for arbitrary edge groups. Finite index here lets us identify the field in the regular representation.

Theorem 6.2. *The canonical map $A_G \to R_G$ identifies R_G with the universal field of fractions of A_G, obtained by inverting the full square matrices over A_G. The affiliated rank of any rectangular A_G-matrix equals its inner rank, the least width of a factorization over A_G.*

Proof. First take a finite graph with edges and use the bipartite subdivision and matrix coordinates of Lemma 5.1. Let $\rho : A_G \to E$ be any unital map to a division ring. It embeds each vertex division ring, and extension to E preserves ranks of matrices over that ring by Gaussian elimination.

For a product XY of the two types of vertex blocks in $M_d(A_G)$, take a minimizing tuple (L_e) from Proposition 5.2, applied in the affiliated representation. Represent each edge space by an idempotent onto L_e over R_{G_e}. Their block diagonal a has the same realization on the two sides by the transport relations. Over E write $\rho(XY) = \rho(Xa)\rho(Y) + \rho(X)\rho((1-a)Y)$. The ranks of the two factors Xa and $(1-a)Y$ are the sums of their local ranks, since the stable-letter diagonals in the right blocks are invertible. The sum is the attained cost. Thus $\operatorname{rk}_E \rho(XY) \leq \operatorname{rk}_G(XY)$, with both sides using expanded matrices. The matrix-unit generation and the stabilization in the proof of Theorem A work equally over the vertex division rings, using inverse stable letters in place of adjoints. Put an arbitrary A_G-matrix in one corner, stabilize it to such a product, and subtract d for each identity block on both sides. This gives $\operatorname{rk}_E \rho(M) \leq \operatorname{rk}_G M$ for every M. For a graph without edges the fundamental ring is the vertex division ring. For an infinite graph, lift the finite expressions

for the entries to a finite connected subgraph containing the relevant tree paths. Restrict ρ to its fundamental ring and use the preservation of affiliated ranks under subgroup inclusion.

By Cohn's theorem, the semifir A_G has a universal field of fractions in which rank is inner rank [5, Theorem 7.5.13 and Corollary 7.5.14]. Apply the inequality with this field as E. Every full square matrix over A_G is invertible in R_G. Hence the map to R_G extends to an embedding of the universal field. Its image is a division ring containing $K[G]$, so it is R_G. This proves the claims. □

The same identification gives $\operatorname{Tor}_j^{A_G}(U_G, U_G) = 0$ for every $j \geq 1$. A universal localization has $\operatorname{Tor}_1^{A_G}(R_G, R_G) = 0$ [17, Lemma 8.6(ii)]. A finitely generated submodule of a finite free module over a semifir is projective, by splitting the surjection onto its first coordinate ideal and inducting on the number of coordinates. Arbitrary submodules of free modules are directed unions of such projectives and are flat, so the higher groups vanish as well. Finally use that U_G is a direct sum of copies of R_G on either side over the division ring. For one amalgam this gives the Tor vanishing in Sánchez-Peralta's Conjecture 7.1 [21] under the finite-index edge and integral vertex-rank hypotheses above.

Let $\Sigma_K(H)$ consist of the square matrices over $K[H]$ that are invertible in U_H.

Corollary 6.3. *If each $K[G_v] \to R_{G_v}$ is the universal localization at $\Sigma_K(G_v)$, then $K[G] \to R_G$ is the universal localization at $\Sigma_K(G)$. In particular this holds for free vertex groups.*

Proof. A unital map $f : K[G] \to S$ inverting $\Sigma_K(G)$ extends uniquely over all vertex division rings, since subgroup inclusion preserves affiliated invertibility. The transported maps agree on the edge division rings because they agree on the edge group rings and preserve inverses. This defines a map on A_G. Each vertex coefficient has a realization $uB^{-1}v$ over its vertex group ring, with B invertible affiliated. Block diagonal and block triangular denominators handle sums and products. Thus any square A_G-matrix M has a realization $M = VB^{-1}W$ over $K[G]$, and the extension of f evaluates it by the same formula. If M is full, Theorem 6.2 makes $\left(\begin{smallmatrix} B & W \\ V & 0 \end{smallmatrix}\right)$ invertible affiliated, by its Schur complement. Its image under f is invertible, so the image of M is invertible. The universal property in Theorem 6.2 now gives exactly the required factorization. For free groups the vertex localization is the usual universal free field, by the Hughes-free identification, Lewin's theorem and the fir property [8, §4.2][12, Theorem 1.1][5, Corollaries 7.5.14 and 7.11.8]. □

Other coefficient fields. The two localization statements also hold for an arbitrary subfield $K \subseteq \mathbb{C}$, with R_H replaced by the affiliated division closure $D_K(H)$ and with the hypothesis that the vertex closures are division rings. Here the analytic formula still uses the complex group von Neumann algebras. In the proof of Lemma 3.2, finite-index coset restriction sends $K[V]$ into $M_s(K[E])$, where $s = [V : E]$. The image of $D_K(V)$ lies in $M_s(D_K(E))$, since a matrix over the division ring $D_K(E)$ invertible in $M_s(U_E)$ is already invertible over $D_K(E)$. The identity column gives the right coset decomposition, and the normal-core grading gives left coordinates and the bimodule coefficient map without taking adjoints in either division ring.

With these coordinates, Lemmas 4.1 and 4.2 apply to right division-ring subspaces. In the upper bound of Proposition 5.2, take the orthogonal projection onto $L_e U_{G_e}$ in the ambient edge algebra. The projection may lie outside the division ring, but its range and quotient compute the same two costs. The rest of simultaneous attainment uses split linear maps and dimensions, so gives integral group-ring ranks by the same stabilization. Every element of $D_K(G)$ is a bordered-matrix realization over $K[G]$, as in the preceding proof. Its integral rank is zero or one, which makes $D_K(G)$ a division ring. The field comparison above uses idempotents on the division subspaces and now applies without a conjugation hypothesis.

6.2. **An exact rank algorithm.**

Theorem 6.4. *There is a uniform algorithm to compute* $\mathrm{rk}_G M$ *for a finite rectangular matrix over* $\mathbb{Q}[G]$ *when the graph is finite and the vertex and edge groups are given by finite free bases. The incidence maps are given by words, with injectivity and finite index at both ends as input hypotheses. Entries of* M *are finite rational sums of words in the vertex generators and the stable letters.*

Proof. Stallings folding decides membership in each incident edge subgroup [10]. After an affirmative test, enumeration of edge words finds a preimage by free reduction. Bass–Serre normal form then decides the word problem in G.

The following two searches apply whenever the input group has a decidable word problem and the rank of M equals its inner rank over an effectively presented ring. Write b for the number of columns, choose a rational $s \geq 1$ bounding the sum of the moduli of the coefficients, and put $P = M^*M/s^2$. The rational numbers

$$\ell_j = \mathrm{Tr}_G(I_b - (I_b - P)^j), \qquad j \geq 1,$$

are computed by polynomial multiplication and the word problem. By the spectral theorem they increase to the rank. For upper bounds, enumerate factorizations $M = BC$ of candidate widths $0, \ldots, \min(a, b)$, where a is the number of rows. In a recursively presented ring, equality is semidecidable by enumerating finite consequences of the relations, so run these tests in parallel. Keep the least certified width u, initially $\min(a, b)$, and stop once $\ell_j > u-1$. An exact-width factorization exists by the inner-rank identity. It is eventually certified, after which a lower bound crosses the stated threshold. This proves termination and correctness without a convergence-rate assumption.

First apply these searches to a free group F itself, taking the factorization ring to be $\mathbb{Q}[F]$. Free reduction decides equality there, and its inner rank agrees with its affiliated rank by Lewin's theorem and Cohn's fir theorem. This local algorithm also decides equality and validity of rational formulas in R_F. Indeed maintain $uB^{-1}v$ representations under sums and products, and before inverting an expression compute its rank as

$$\mathrm{rk}_F \begin{pmatrix} B & v \\ u & 0 \end{pmatrix} - \mathrm{size}(B).$$

Discard inversion of zero. The inverse of a nonzero expression is represented by the same bordered denominator, with the appropriate row and column selectors and a minus sign. Equality reduces to testing the difference.

At every vertex and edge, valid formulas give enumerable names and effective arithmetic for its division ring. Edge maps are computed by substitution of the specified basis words, since the affiliated subgroup embedding preserves their nonzero denominators. Present A_G by the vertex arithmetic relations, stable-letter inverses and the transport relations on these enumerable names. The resulting recursive presentation and the word-problem algorithm fulfill the two search requirements for G, by Theorem 6.2. □

Example. Use two cyclic vertices with two parallel edges, so $G = \langle a, b, t \mid a^2 = b^3,\ tb^6t^{-1} = a^6 \rangle$. The edge indices are $(2, 3)$ and $(6, 6)$. The relations reduce

$$\begin{pmatrix} t - a^6t + tb^6 & t(1+b) \\ at & -a^2 + b^3 + at(1+b) \end{pmatrix} = M = \begin{pmatrix} t \\ at \end{pmatrix} \begin{pmatrix} 1 & 1+b \end{pmatrix}.$$

This width-one certificate is paired with the lower bound $\ell_3 = 263/648 > 0$ for $P = M^*M/36$ after the cancellations. Indeed the first three unnormalized moments of M^*M are $6, 44, 360$, by taking the identity coefficients of the corresponding powers of $2(3 + b + b^{-1})$. Thus the two bounds certify rank one. Here the lower bound can also be read off from the unit in the upper-left corner, and $\ell_1 = \mathrm{tr}(M^*M)/36 = 1/6 > 0$ already certifies rank at least one,

as for every nonzero matrix. Trace moments become informative only for certificates of rank greater than one.

Effective transfer. The two rank searches above apply over any named coefficient field $K \subseteq \mathbb{C}$ under the hypotheses of Theorem 6.2. The required data are a word-problem algorithm for G, complex approximations to named coefficients to arbitrary precision, and a recursively enumerable presentation of A_G with computable names for the coefficients and group generators. Approximate the j-th trace polynomial from below within 2^{-j} and keep the running maximum. The factorization search uses the ring presentation, and the inner-rank identity guarantees that it eventually certifies the exact integer. The same argument works relative to an oracle supplying the required data.

Corollary 6.5 (One vertex rank oracle)**.** *Let G be the fundamental group of a finite connected graph of groups with finite-index edge inclusions and torsion-free vertex groups satisfying the strong Atiyah conjecture over $\mathbb{Q}$. The input gives finite presentations for the vertices and a finite generating list for each edge, with its two lists of images in the incident vertex generators. Injectivity and finite index are input promises. For any specified vertex v_0, the exact rank problems for matrices over $\mathbb{Q}[G]$ and $\mathbb{Q}[G_{v_0}]$ are Turing-equivalent, uniformly in these data.*

Proof. At an incidence $H \leq V$, enumerate finite transitive right permutation actions of the presented vertex group V with a marked point fixed by the listed generators of H. For each candidate choose words s_i carrying the marked point to i, with $s_i = 1$ at the marked point. Its stabilizer is generated by the Schreier words $s_i x s_{ix}^{-1}$ for signed vertex generators x. Search in parallel for proofs from the vertex relators that each Schreier word equals a word in the listed edge images. A successful search proves that the stabilizer is H, since it already contains H. The true finite coset action has such witnesses and will be found. Its table tests membership, and tracking the Schreier words along a path rewrites any accepted word in the edge generators.

An edge matrix has the same rank in either vertex. For the opposite transfer, restrict a matrix M over $\mathbb{Q}[V]$ to right $\mathbb{Q}[H]$-modules with basis $r_i = s_i^{-1}$. The table locates i for each product $gr_j = r_i h$ by reading $(gr_j)^{-1}$ from the marked point. It then rewrites $h = r_i^{-1} gr_j$ in the edge generators. The resulting matrix M_{res} satisfies

$$\operatorname{rk}_V M = [V : H]^{-1} \operatorname{rk}_H M_{\mathrm{res}}$$

by Lemma 3.2. Following spanning-tree paths therefore propagates the oracle at v_0 to every vertex.

Testing the rank of $1 - g$ decides the vertex word problems. The incidence tables then give the word problem in G by Bass–Serre normal form. The bordered-matrix test above decides validity and equality of rational formulas over the vertex and edge group rings. Their arithmetic and the stable-letter and transport relations present A_G relative to the oracle. Effective transfer gives the rank algorithm for G relative to that oracle. The converse reduction is the word inclusion of G_{v_0} in G, which preserves affiliated ranks. □

6.3. **Changing coefficients and detecting point eigenvalues.** The dimension of an eigenspace bounds the number of its algebraic conjugates once we know that conjugating coefficients preserves ranks. A field automorphism of $\mathbb{C}$ need not preserve the involution. The universal field lets us work with ordinary ring operations. Proposition 6.6 and Corollary 6.7 below are cases of the independence, strong algebraic eigenvalue and center conjectures, which Jaikin-Zapirain proved for sofic groups [9, Corollaries 1.5–1.7].

Proposition 6.6. *Suppose the vertex ranks are integral over $\mathbb{C}$ and the edge inclusions have finite index at both ends. If $\sigma \in \mathrm{Aut}(\mathbb{C})$ preserves the ranks of all complex group-ring matrices at one vertex of the connected graph, then it extends uniquely to an automorphism of R_G fixing G, and it preserves matrix ranks.*

Proof. At an inclusion $E \leq V$ of finite index, the property passes down by subgroup inclusion. To pass up, express a vertex matrix in coset coordinates over $\mathbb{C}[E]$. Changing coefficients before or after writing these coordinates gives the same matrix, and restriction multiplies ranks by $[V : E]$ on both sides (Lemma 3.2). Along a finite path the property therefore reaches every vertex.

At a vertex, represent a division-ring coefficient by $uB^{-1}v$ over its group ring. After applying σ the denominator still has full rank. The difference of two such expressions has a block diagonal denominator, and its value is zero exactly when the bordered-matrix rank equals the size of that denominator, as in Theorem 6.4. Consequently changing coefficients gives a well-defined ring map on the vertex division ring. Its inverse comes from σ^{-1}. The vertex maps agree on the transported edge division rings, since they agree on the group rings which generate them by division. Fixing the stable letters now defines an automorphism of A_G. It preserves full matrices because it preserves factorizations and their widths in both directions. Theorem 6.2 extends it to R_G, where Gaussian elimination computes affiliated ranks. Division generation gives uniqueness. The same presentation and argument apply to an infinite graph, since every word and every matrix has finite support. □

Corollary 6.7. *Suppose the proposition applies to every $\sigma \in \mathrm{Aut}(\mathbb{C})$, as it does for a finite-index graph of free groups. Let $K \subseteq \mathbb{C}$ be a subfield and $T \in M_n(K[G])$. Every point eigenvalue λ of T on $\ell^2(G)^n$ is algebraic over K, and*

$$[K(\lambda) : K] \dim_G \ker(T - \lambda I) \leq n.$$

The same bound holds for a singular pencil $T - \lambda I$ with $T \in M_n(D_K(G))$. In addition $D_K(G) \cap \mathbb{C} = K$.

Proof. Free groups are sofic, so coefficient automorphisms preserve their ranks by Jaikin-Zapirain's independence theorem [9, Corollary 1.7]. In general fix σ acting trivially on K. Its extension fixes $D_K(G)$ pointwise and carries the right kernel of $T - \lambda I$ over R_G to the kernel at $\sigma(\lambda)$. Each has dimension $r = \dim_G \ker(T - \lambda I) > 0$. Eigenspaces for distinct central scalars sum directly: applying $\prod_{\mu \neq \nu}(T - \mu I)$ to a finite relation isolates its term at ν. Thus every finite subset of the orbit of λ has at most n/r elements. A transcendental λ has infinite orbit, since one can include it in a transcendence basis, translate it by integers, and extend the automorphisms to the algebraic closure. An algebraic λ has exactly $[K(\lambda) : K]$ conjugates, each realized by extension to $\mathbb{C}$. These facts give the bound and show that the fixed field of $\mathrm{Aut}(\mathbb{C}/K)$ is K, proving the intersection assertion. □

There is also a bound on the number of point eigenvalues when torsion is present. If $d(G) < \infty$ and $\Lambda_{\mathbb{C}}(G) = d(G)^{-1}\mathbb{Z}$, an n-by-n complex group-ring matrix has at most $nd(G)$ point eigenvalues. The same polynomial separation works on the right kernels over $\mathcal{U}(G)$. Their dimensions add to at most n, and every nonzero kernel has dimension at least $1/d(G)$.

7. Tree actions and the rational matrix groups

Corollary 7.1. *Suppose G acts coordinatewise on a product of finitely many locally finite trees, with finite stabilizers of vertex tuples and $d(G) < \infty$. Then $\Lambda_{\mathbb{C}}(G) = d(G)^{-1}\mathbb{Z}$.*

Proof. Subdivide to remove inversions. A stabilizer of a tuple involving midpoints has a subgroup of finite index fixing chosen endpoints of the corresponding edges, hence is still finite. Act on the first tree. The stabilizer of an incident edge has finite index in the vertex stabilizer since its index is the size of an orbit in a finite set of incident edges. Each vertex stabilizer acts on the remaining factors and has finite tuple stabilizers there. Its torsion denominator divides $d(G)$. Induction, starting with finite groups acting on a point, and Theorem A give the result. For a finite group F, von Neumann rank is the ordinary complex rank of the finite regular-representation matrix divided by $|F|$. □

For torsion-free groups in Corollary 7.1, the same induction gives the eigenvalue conclusion of Corollary 6.7. Start with the trivial group at a vertex tuple. At each tree, Proposition 6.6 passes coefficient invariance from a vertex stabilizer to the acting group, for each automorphism of $\mathbb{C}$.

7.1. **Finitely presented simple groups.** For even integers $n \geq 218$ and $m \geq 350$, Burger and Mozes construct torsion-free finitely presented infinite simple lattices $\Gamma_{n,m} \leq \mathrm{Aut}(T_n) \times \mathrm{Aut}(T_m)$, where T_d is the regular tree of degree d. The two tree actions give amalgams of finitely generated free groups. The edge group has index n in each of the two vertex groups for the T_n-action, and index m in each for the T_m-action [3, the theorem and Remark (3) in §0, and §1]. See also [4, Theorem 5.5].

Rattaggi gives an explicit example Σ_0 inside the $(12, 8)$-lattice Σ. The latter has the ten generators $a_1, \ldots, a_6, b_1, \ldots, b_4$ and the twenty-four square relators in [19, Example 3.4]. The index-four subgroup Σ_0 is the kernel of the map to $(\mathbb{Z}/2\mathbb{Z})^2$ sending each a_i to $(1, 0)$ and each b_j to $(0, 1)$. It is an infinite, finitely presented, torsion-free simple group [19, §1 and Theorem 3.5]. Writing F_r for the free group of rank r, its two decompositions are

$$F_7 *_{F_{73}} F_7 \cong \Sigma_0 \cong F_{11} *_{F_{81}} F_{11}.$$

The two edge inclusions have index 12 in the first amalgam and index 8 in the second [18, Proposition 1.4 and Theorem 2.47(4)].

Let G be any of these simple groups. A vertex-tuple stabilizer is the intersection of the discrete lattice with a compact subgroup of the ambient automorphism group. It is finite and hence trivial by torsion-freeness. Corollary 7.1 gives $\Lambda_{\mathbb{C}}(G) = \mathbb{Z}$. For $0 \neq x \in \mathbb{C}[G]$, the vector $x\delta_1$ is nonzero. Faithfulness of the trace gives $0 < \mathrm{rk}_G\, x \leq 1$, so $\mathrm{rk}_G\, x = 1$ and x acts injectively on $\ell^2(G)$. In particular $\mathbb{C}[G]$ is a domain, since $xy = 0$ implies $x(y\delta_1) = 0$. An idempotent e satisfies $e(1 - e) = 0$, so e is 0 or 1. Both algebraic conclusions hold over any characteristic-zero field, since the finitely many coefficients of a putative counterexample generate a subfield that embeds in $\mathbb{C}$. The integer rank statement also holds over every subfield of $\mathbb{C}$. For a finite CW complex with fundamental group G, its L^2-Betti numbers are integers, since each is the number of cells in the corresponding degree minus the ranks of the two adjacent cellular boundary matrices of the universal cover. For the group itself one can take its finite classifying space, the quotient of the product of trees.
These implications of complex rank integrality apply to any torsion-free group satisfying the complex conjecture. Coefficients of positive characteristic raise different questions. In earlier work we constructed torsion-free hyperbolic groups with nontrivial idempotents in positive characteristic [16], and diffuse groups whose group algebras in a prescribed positive characteristic are domains admitting no division-ring embedding [15].

7.2. **Subgroups of the rational matrix groups.**

Theorem 7.2. *Every subgroup $H \leq \mathrm{GL}_2(\mathbb{Q})$ satisfies the strong Atiyah conjecture over $\mathbb{C}$ with denominator $d(H)$. One has $d(\mathrm{GL}_2(\mathbb{Q})) = 24$ and $d(\mathrm{SL}_2(\mathbb{Q})) = 12$. The same conclusions about a given group's denominator hold over every subfield of $\mathbb{C}$.*

Proof. First consider groups with determinant in $\{1, -1\}$. Every finite subgroup F of $\mathrm{GL}_2(\mathbb{Q})$ preserves the full lattice $\sum_{g \in F} g\mathbb{Z}^2$ and hence is rationally conjugate into $\mathrm{GL}_2(\mathbb{Z})$. A finite such group preserves a positive definite real inner product. Its rotation subgroup is finite cyclic, and its index is at most two. The traces of its rotations are rational algebraic integers in $[-2, 2]$, so the possible rotation orders are $1, 2, 3, 4, 6$. Thus $d(\mathrm{GL}_2(\mathbb{Q})) \mid 24$ and $d(\mathrm{SL}_2(\mathbb{Q})) \mid 12$.

For a finite set of primes S put $A_S = \mathbb{Z}[1/\prod_{q \in S} q]$. We prove the assertion for all subgroups of $\mathrm{GL}_2(A_S)$ with determinant ± 1, by induction on $|S|$. At $S = \varnothing$ use the dual tree of the Farey tessellation [22]. The action of $\mathrm{GL}_2(\mathbb{Z})$ permutes its ideal triangles, and their stabilizers are finite: a projective transformation fixing three distinct rational slopes

is scalar, with scalar kernel $\{\pm I\}$ here. Subdivision removes inversions and still has finite stabilizers. Corollary 7.1 applies to each subgroup.

Choose $q \in S$ and use the Bruhat–Tits tree of $\mathrm{PGL}_2(\mathbb{Q}_q)$, whose vertices are homothety classes of $\mathbb{Z}_q$-lattices. An element of determinant ± 1 fixing such a class fixes the lattice itself, by comparing determinant valuations. It also preserves the bipartite type, so there are no edge inversions. A vertex can be represented by $g\mathbb{Z}_q^2$ with $g \in \mathrm{GL}_2(\mathbb{Z}[1/q])$. Indeed, after rescaling, the lattice lies between $q^m\mathbb{Z}_q^2$ and $\mathbb{Z}_q^2$. Take the inverse image in $\mathbb{Z}^2$ of its residue subgroup modulo q^m. This is a full lattice of q-power index. A matrix of its integral basis has determinant $\pm q^k$ and its columns generate the prescribed q-adic lattice. A vertex stabilizer in our subgroup conjugated by g^{-1} is integral at q and still integral outside S. It therefore has entries in $A_{S\setminus\{q\}}$ and determinant ± 1. Each edge inclusion has finite index by local finiteness. The induction hypothesis and Theorem A prove the assertion for this whole subgroup, with its own denominator. Every finitely generated determinant-± 1 subgroup is covered by some S. Finite support and subgroup rank preservation finish this case.

For general H first assume that the image of $|\det| : H \to \mathbb{Q}_{>0}^{\times}$ is finitely generated. It is free abelian. If its rank is positive, project onto one $\mathbb{Z}$-coordinate and act by translations on a line. The vertex and edge groups are all the kernel, and every finite subgroup lies in that kernel. Induction on the determinant rank using Theorem A proves the assertion. In an arbitrary H each matrix uses a finitely generated subgroup, whose denominator divides $d(H)$. Finite-subgroup averages give the reverse inclusion of rank groups over any coefficient subfield, since these averages have rational coefficients.

To see the exact denominators take

$$a = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \qquad b = \begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix}, \qquad s = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

The orders of a, b are $4, 6$. Conjugation by s inverts each, giving dihedral subgroups of orders $8, 12$. The lower bounds follow. Because a, b already occur in $\mathrm{SL}_2(\mathbb{Z})$, the determinant-one groups in Theorem B all have denominator 12. □

Proposition 7.3. *There is a selfadjoint element of* $\mathbb{Z}[\mathrm{SL}_2(\mathbb{Z})]$ *of kernel dimension* $1/12$ *and one of* $\mathbb{Z}[\mathrm{GL}_2(\mathbb{Z})]$ *of kernel dimension* $1/24$*. Each has the same kernel dimension in every overgroup.*

Proof. In the determinant-one case put $A = \langle a \rangle$, $B = \langle b \rangle$. In the general linear case put $A = \langle a, s \rangle$, $B = \langle b, s \rangle$. Write p, q for their respective averaging projections. In both cases there is a subgroup E of index two in A and contained in B: take $\langle -I \rangle$ or $\langle -I, s \rangle$, respectively. On the B-fixed space $q\ell^2$, the operator p acts as $(1 + a)/2$, by the cosets of E in A. A vector ξ there annihilated by p satisfies $a\xi = -\xi$ and $b\xi = \xi$. Its pointwise absolute value is invariant under $a^{-1}b = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ of infinite order. Square summability forces $\xi = 0$.

It follows by taking the adjoint that $qp : p\ell^2 \to q\ell^2$ has dense image. The positive operator $1 - p + q$ has kernel $p\ell^2 \cap \ker q$, with dimension $\tau(p) - \tau(q) = 1/|A| - 1/|B|$. This is $1/12$ or $1/24$, respectively. Multiplying $1 - p + q$ by 12 or 24 clears denominators. Rank preservation under group inclusion gives the assertion about overgroups. This also completes the proof of Theorem B. □

Proposition 7.4. *The minimum index of a torsion-free finite-index subgroup of* $\mathrm{SL}_2(\mathbb{Z}[1/2])$ *is* 24*, and every such index is a multiple of* 24*.*

Proof. On n cosets of such a subgroup both cyclic groups $\langle a \rangle, \langle b \rangle$ act freely, hence $n = 12k$. The permutation a is a product of $3k$ four-cycles, of sign $(-1)^k$. The permutation b is a product of $2k$ six-cycles, of sign $+1$. Yet $a^{-1}b$ is the square of the unipotent with upper right entry $1/2$, so its sign is $+1$. Thus k is even.

Reduction modulo 3 maps onto $\mathrm{SL}_2(\mathbb{F}_3)$ of order 24, since the elementary generators lift. A finite-order matrix in its kernel different from I would have trace -1: among

$\{-2,-1,0,1,2\}$ only $-1,2$ reduce to 2, and a finite-order matrix of trace 2 is I. Writing it as $I+3C$ gives $\operatorname{Tr} C=-1$ and $\det C=1/3$ from the determinant equation. This is impossible over $\mathbb{Z}[1/2]$. Its kernel realizes the claimed index. □

8. Quaternion algebras over global fields

The tree argument also applies when a single noncompact archimedean group remains after the nonarchimedean factors have been removed. A global function field here has characteristic $p>0$, while the group ring in the conclusion has characteristic zero.

Corollary 8.1. *Let k be a global field and D a quaternion algebra over k, split at at most one archimedean place. Write $D^1(k)$ for its reduced norm-one group. Every subgroup $H\le D^1(k)$ with $d(H)<\infty$ satisfies the strong Atiyah conjecture over $\mathbb{C}$ with denominator $d(H)$. The boundedness hypothesis holds for every such H when k is a number field, and also when k is a function field and D is a division algebra. In particular every subgroup of $\mathrm{SL}_2(k)$ is covered for k imaginary quadratic.*

Proof. Fix a finitely generated subgroup $J\le H$. Choose a k-basis of D containing 1 and a nonempty finite set S of places containing all archimedean places and all denominator places of the multiplication table and the coordinates of chosen generators of J and their inverses. The elements of J then have S-integral coordinates. The diagonal copy of J in $\prod_{v\in S} D^1(k_v)$ is discrete. Indeed an S-integer of sufficiently small absolute value at every place in S is zero by the product formula, applied one coordinate at a time.

At a nonsplit place the local norm-one group is compact. At a split nonarchimedean place it is $\mathrm{SL}_2(k_v)$ and acts on a locally finite Bruhat–Tits tree with compact vertex stabilizers [22, 23]. Following the proof of Corollary 7.1, act on these trees one at a time. The last stabilizers are discrete subgroups Δ of a product of compact factors and at most one factor $\mathrm{SL}_2(\mathbb{R})$ or $\mathrm{SL}_2(\mathbb{C})$. With just compact factors they are finite. Otherwise projection into $\mathrm{Isom}(\mathbb{H}^3)$, via $\mathrm{SL}_2(\mathbb{R})\subset\mathrm{SL}_2(\mathbb{C})$ for a real factor, has discrete image and finite kernel: the archimedean factor acts through a closed isometry subgroup with finite kernel, and the inverse image of a compact set in that subgroup, with the previous compact factors included, is compact. Its intersection with the closed discrete group Δ is finite.

In the archimedean case k is a number field, so left multiplication on D, followed by restriction to $\mathbb{Q}$, gives a faithful rational linear representation. The finitely generated J has a torsion-free normal subgroup J_0 of finite index by Selberg's lemma (one can reduce its rational matrices modulo an odd prime outside the denominators). Thus $\Delta\cap J_0$ injects as a torsion-free discrete hyperbolic isometry group. It acts freely on $\mathbb{H}^3$ and is the fundamental group of its quotient three-manifold. By Kielak–Linton [11, Theorem 1.4] it belongs to Linnell's class $\mathcal{C}$. This class is closed under extensions with finite quotient. Hence $\Delta\in\mathcal{C}$, and Linnell's bounded-torsion theorem [13, Theorem 1.5] gives the exact conjectured grid for Δ. Applying Theorem A up the trees proves the conjecture for J, and finite supports give it for H.

For the two boundedness assertions, any finite subgroup in the rational linear representation preserves a full integral lattice. In fixed degree finite subgroups of $\mathrm{GL}_m(\mathbb{Z})$ inject modulo 3. To see this, take a putative prime-order element $I+3^aB$ in the kernel, $a\ge 1$, with $B\not\equiv 0 \pmod 3$. For prime order $\ell\neq 3$ the leading binomial term survives modulo 3^{a+1}; for $\ell=3$ it survives modulo 3^{a+2}. This gives a uniform finite subgroup bound.

In the function-field division case let the exact constant field be $\mathbb{F}_q$. The $\mathbb{F}_q$-span of a finite subgroup F in D is a finite domain and hence a finite field. Its compositum with k is a commutative subfield of the quaternion division algebra, so has degree at most two over k. A constant extension has its full constant-field degree since k has exact constant field $\mathbb{F}_q$. Thus either F is central of order dividing $\gcd(2,q-1)$, or it lies in $\mathbb{F}_{q^2}^{\times}$ with reduced norm $x^{q+1}=1$. In both cases $|F|\mid q+1$. For the split group over a function field, upper

triangular unipotents from n-dimensional additive $\mathbb{F}_p$-subspaces have order p^n, accounting for the boundedness qualification. □

Appendix A. A rotation proof for group amalgams

Let $G = A *_E B$. We give a proof of Theorem 2.1 in this case using the normal form for amalgams. Form $\widehat{G}$ by amalgamating countably many labelled copies A_i of A and B_j of B, all over the same E. Each pair A_i, B_j generates a copy of G. A subscript on a vertex matrix denotes its image in the indicated copy. Unless a ring is specified, kernels and images are taken on columns over $\mathcal{U}(\widehat{G})$.

Lemma A.1 (Locality under permutations)**.** *In the subgroup generated by A_i and all the B_j, the subalgebra of the group von Neumann algebra fixed by all finite permutations of the B_j is $\mathcal{N}(A_i)$. The analogous assertion holds with A and B interchanged. In $\mathcal{N}(\widehat{G})$ the subalgebra fixed by both families of finite permutations is $\mathcal{N}(E)$. These assertions hold entrywise for matrices.*

Proof. The sequence of vertex labels in a reduced word is independent of the choice of its syllable representatives. Every element of the indicated subgroup outside A_i has a B-label in that sequence. Exchanging this label with infinitely many labels absent from the word produces infinitely many distinct elements. In $\widehat{G}$, every element outside E has at least one label, which can likewise be exchanged within its family.

For a bounded invariant operator x, its Fourier coefficients are the coordinates of $x\delta_1$ in the group basis. They are square-summable and constant on each orbit, hence zero on every infinite orbit. Thus they are supported on A_i in the first assertion, and on E in the last. Conditional expectation onto the subgroup algebra preserves these coefficients and kills the other coefficients. It therefore fixes x, by faithfulness of the group trace. Conversely, each stated subgroup is fixed pointwise. The same proof applies to every matrix entry. □

Lemma A.2 (A constant-rank pencil)**.** *Let Z_0, Z_1 be bounded rectangular matrices of the same size over a finite tracial von Neumann algebra. If $\operatorname{rk}(Z_0 + tZ_1) = \operatorname{rk} Z_0$ for real $t \neq 0$ arbitrarily close to zero, then $Z_1(\ker Z_0) \subseteq \operatorname{im} Z_0$ over the affiliated ring.*

Proof. Let p be the range projection of Z_0 and set $q = 1 - p$. For $t \neq 0$, left multiplication by the invertible $p + t^{-1}q$ preserves rank and gives

$$(p + t^{-1}q)(Z_0 + tZ_1) = Z_0 + qZ_1 + tpZ_1.$$

Rank is lower semicontinuous in operator norm. Indeed, if $J_k \to J$ and $\|J_k - J\| < \eta$, the range of the source spectral projection $1_{[\eta,\infty)}(|J|)$ has zero intersection with $\ker J_k$, since J is bounded below by η there. Its dimension is at most $\operatorname{rk} J_k$. Letting η decrease to zero gives the assertion.

Consequently $\operatorname{rk}(Z_0 + qZ_1) \leq \operatorname{rk} Z_0$. Projection by p takes $\operatorname{im}(Z_0 + qZ_1)$ onto $\operatorname{im} Z_0$, and its kernel on this image is $qZ_1(\ker Z_0)$. Thus

$$\operatorname{rk}(Z_0 + qZ_1) = \operatorname{rk} Z_0 + \dim\big(qZ_1(\ker Z_0)\big).$$

The last module is the image of a finite matrix, so dimension zero forces it to vanish. This gives the required inclusion. □

Lemma A.3 (Rotating two copies)**.** *Let $a \in M_m(\mathcal{N}(A))$, $b \in M_m(\mathcal{N}(B))$ and $R_{ij} = a_i + b_j$. If $i \neq k$ and $j \neq l$, then*

$$R_{ij}(\ker R_{kj}) \subseteq \operatorname{im} R_{il}. \tag{A.1}$$

Proof. For a scalar 2-by-2 unitary V, represent G in $M_2(\mathcal{N}(\widehat{G}))$ by

$$x \in A \longmapsto \operatorname{diag}(x_1, x_2), \qquad y \in B \longmapsto V \operatorname{diag}(y_1, y_2)V^*.$$

The two maps agree on E. The normalized matrix trace of a represented group element is its group trace. Indeed, expanding a diagonal entry for a reduced alternating word of positive syllable length gives a sum of scalar multiples of words in labelled copies. The syllables still alternate between types A and B, and all lie outside E. Every term therefore has trace zero by normal form. A nonidentity element of E has trace zero in both diagonal blocks. The identity has normalized trace one.

Trace preservation identifies the trace GNS space of the image group algebra with $\ell^2(G)$, intertwining their left multiplication representations. The resulting map extends to a faithful normal trace-preserving homomorphism from $\mathcal{N}(G)$ onto the von Neumann algebra generated by the image. Under matrix amplification it preserves ranks computed with the normalized trace on the two-copy factor. Its restrictions to the two vertex von Neumann algebras are given by the same diagonal and conjugated-diagonal formulas.

For real t, take $V = (1+t^2)^{-1/2} \left(\begin{smallmatrix} 1 & t \\ -t & 1 \end{smallmatrix}\right)$. Multiply the represented $a+b$ on the right by V, amplified by I_m, and remove the nonzero scalar factor. The resulting pencil is

$$Z(t) = \begin{pmatrix} R_{11} & tR_{12} \\ -tR_{21} & R_{22} \end{pmatrix}.$$

It has unnormalized rank $2\,\mathrm{rk}_G(a+b)$ for every t. For $R_{22}x = 0$, the linear term sends $(0,x)$ to $(R_{12}x, 0)$. Lemma A.2 puts this vector in $\operatorname{im} Z(0)$, hence $R_{12}x \in \operatorname{im} R_{11}$. Relabelling the copies gives (A.1). □

Proof of Theorem 2.1 for $A *_E B$. Take $X \in M_{a,m}(\mathcal{U}(A))$ and $Y \in M_{m,b}(\mathcal{U}(B))$. Write $K_i = \ker X_i$, $I_j = \operatorname{im} Y_j$, and $H_{ij} = K_i \cap I_j$. Restricting X_i to I_j and using the trace-preserving subgroup embedding gives

$$h := \dim H_{ij} = \mathrm{rk}_B\, Y - \mathrm{rk}_G(XY)$$

for every pair. Take the projection joins

$$L_i = \bigvee_j H_{ij}, \qquad M_j = \bigvee_i H_{ij},$$

and denote their projections by p_i, q_j. Each H_{ij} is obtained by taking the meet of projections in the algebra generated by A_i, B_j. Hence p_i lies in the matrix algebra over the subgroup generated by A_i and all the B_j. Finite permutations of the B_j fix p_i, so Lemma A.1 gives $p_i \in M_m(\mathcal{N}(A_i))$. A permutation taking A_i to A_k takes each H_{ij} to H_{kj} and hence p_i to p_k. These projections are therefore copies of one matrix $p \in M_m(\mathcal{N}(A))$. The same argument gives q_j as copies of a matrix $q \in M_m(\mathcal{N}(B))$. We have

$$L_i \subseteq K_i, \qquad M_j \subseteq I_j, \qquad L_i \cap M_j = H_{ij}.$$

Fix i, l. We show that $L_i + M_l$ contains every H_{kj}. This is immediate if $k = i$ or $j = l$. Otherwise, for $z \in H_{kj}$, both p_k and q_j fix z. Apply (A.1) to $p, -q$, obtaining

$$p_i z - z = (p_i - q_j)z \in \operatorname{im}(p_i - q_l) \subseteq L_i + M_l.$$

Since $p_i z \in L_i$, also $z \in L_i + M_l$. The finite sum is a projection-generated module, and both summands are contained in the join of all the H_{kj}. It follows that

$$W := \bigvee_{k,j} H_{kj} = L_i + M_l$$

for every i, l. Its projection w is fixed by both permutation families, so Lemma A.1 places w in $M_m(\mathcal{N}(E))$.

For the pair i, j, write $W = L_i + M_j$. If $u + v \in I_j$ with $u \in L_i$ and $v \in M_j$, then $u \in K_i \cap I_j = H_{ij} \subseteq M_j$. Hence $W \cap I_j = M_j$. Also $X_i W = X_i M_j$, and the restriction to M_j has kernel H_{ij}. We obtain

$$\mathrm{rk}_A(Xw) = \dim M_j - h, \qquad \mathrm{rk}_B((1-w)Y) = \mathrm{rk}_B\, Y - \dim M_j.$$

The sum is $\mathrm{rk}_G(XY)$. Every other edge projection gives an upper bound by the splitting $XY = XwY + X(1-w)Y$, proving (2.1).

Let W' be the range of any minimizing edge projection. Its two costs are the same for every pair i, j. Computing them by restriction gives

$$h = \dim(W' \cap K_i) + \dim(W' \cap I_j) - \dim W' \;\le\; \dim(W' \cap H_{ij}) \;\le\; h.$$

The first inequality is the dimension formula inside W'. Faithfulness gives $H_{ij} \subseteq W'$ for every pair, hence $W \subseteq W'$.

Finally let r be the projection of H_{11} and $\mathbb{E}_m$ the entrywise expectation from the $(1,1)$ copy of G onto E. An edge projection containing r contains each of its translates under relabelling, so w is the least edge projection containing r. For an edge projection f, faithfulness and bimodularity give

$$r \le f \quad\Longleftrightarrow\quad (1-f)\mathbb{E}_m(r)(1-f) = 0 \quad\Longleftrightarrow\quad \operatorname{supp} \mathbb{E}_m(r) \le f.$$

Thus $w = \operatorname{supp} \mathbb{E}_m(r)$, the claimed canonical minimizer. □

References

[1] S. T. Belinschi, H. Bercovici and W. Liu, *The atoms of operator-valued free convolutions*, J. Operator Theory **85** (2021), 303–320. doi:10.7900/jot.2019dec07.2283. The theorem numbering used here is that of arXiv:1903.09002v2.

[2] S. K. Berberian, *The maximal ring of quotients of a finite von Neumann algebra*, Rocky Mountain J. Math. **12** (1982), 149–164.

[3] M. Burger and S. Mozes, *Finitely presented simple groups and products of trees*, C. R. Acad. Sci. Paris Sér. I Math. **324** (1997), 747–752.

[4] M. Burger and S. Mozes, *Lattices in product of trees*, Publ. Math. Inst. Hautes Études Sci. **92** (2000), 151–194. doi:10.1007/BF02698916.

[5] P. M. Cohn, *Free ideal rings and localization in general rings*, New Mathematical Monographs, vol. 3, Cambridge University Press, Cambridge, 2006. doi:10.1017/CBO9780511542794.

[6] S. Douglas and N. Mghirbi, *Integral coefficients suffice for the universal strong Atiyah conjecture*, companion manuscript, 2026.

[7] S. P. Fisher and P. Sánchez-Peralta, *Division rings for group algebras of virtually compact special groups and 3-manifold groups*, J. Comb. Algebra **10** (2026), 153–193. doi:10.4171/JCA/89.

[8] F. Henneke and D. López-Álvarez, *Pseudo-Sylvester domains and skew Laurent polynomials over firs*, J. Algebra Appl. **21** (2022), article 2250168. arXiv:2006.08454.

[9] A. Jaikin-Zapirain, *The base change in the Atiyah and the Lück approximation conjectures*, Geom. Funct. Anal. **29** (2019), 464–538.

[10] I. Kapovich and A. Myasnikov, *Stallings foldings and subgroups of free groups*, J. Algebra **248** (2002), 608–668. doi:10.1006/jabr.2001.9033.

[11] D. Kielak and M. Linton, *Group rings of three-manifold groups*, Proc. Amer. Math. Soc. **152** (2024), 1939–1946. doi:10.1090/proc/16716. See also arXiv:2303.15907v3, Corollary 1.5(1).

[12] J. Lewin, *Fields of fractions for group algebras of free groups*, Trans. Amer. Math. Soc. **192** (1974), 339–346. doi:10.1090/S0002-9947-1974-0338055-4.

[13] P. A. Linnell, *Division rings and group von Neumann algebras*, Forum Math. **5** (1993), 561–576. doi:10.1515/form.1993.5.561.

[14] W. Lück, *L^2-invariants: theory and applications to geometry and K-theory*, Ergebnisse der Mathematik, 3rd series, vol. 44, Springer, Berlin, 2002.

[15] N. Mghirbi and S. Douglas, *Group algebras that are domains but embed in no division ring*, 2026. `https://finite.works/papers/domain-nonembedding/`.

[16] N. Mghirbi and S. Douglas, *Kaplansky's conjectures fail for torsion-free hyperbolic groups in positive characteristic*, 2026. arXiv:2611.04599.

[17] A. Neeman and A. Ranicki, *Noncommutative localisation in algebraic K-theory I*, Geom. Topol. **8** (2004), 1385–1425. doi:10.2140/gt.2004.8.1385.

[18] D. Rattaggi, *Computations in groups acting on a product of trees: normal subgroup structures and quaternion lattices*, doctoral thesis no. 15538, ETH Zürich, 2004. doi:10.3929/ethz-a-004794219.

[19] D. Rattaggi, *A finitely presented torsion-free simple group*, J. Group Theory **10** (2007), 363–371. doi:10.1515/JGT.2007.028.

[20] H. Reich, *On the K- and L-theory of the algebra of operators affiliated to a finite von Neumann algebra*, K-Theory **24** (2001), 303–326.

[21] P. Sánchez-Peralta, *Universal localizations, Atiyah conjectures and graphs of groups*, Geom. Funct. Anal. **35** (2025), 842–876. doi:10.1007/s00039-025-00710-4.
[22] J.-P. Serre, *Trees*, translated by J. Stillwell, Springer, Berlin, 1980, corrected second printing, 2003.
[23] J. Voight, *Quaternion algebras*, Graduate Texts in Mathematics, vol. 288, Springer, Cham, 2021. doi:10.1007/978-3-030-56694-4.